Computer Science Department

TECHNICAL REPORT

'DREAMWORLD': THE USE OF COMPUTERIZED ELECTRONIC MULTIMEDIA DISPLAY ENVIRONMENTS IN EDUCATION AND ART

By Jacob T. Schwartz

Technical Report #211
March 1986

NYU COMPSCI TR-211
Schwartz, Jacob T
'Dreamworld' c.2

NEW YORK UNIVERSITY

Department of Computer Science
Courant Institute of Mathematical Sciences
251 MERCER STREET, NEW YORK, N.Y. 10012

'DREAMWORLD': THE USE OF COMPUTERIZED ELECTRONIC MULTIMEDIA DISPLAY ENVIRONMENTS IN EDUCATION AND ART

By Jacob T. Schwartz

Technical Report #211
March 1986

'Dreamworld':
The Use of Computerized Electronic Multimedia Display Environments in Education and Art

Jacob T. Schwartz
March 25, 1986

Contents:

1. Introduction
2. Technologies
3. The Computerized Electronic Environment as a Medium
4. Educational Use of the Multimedia Environment
5. Some Additional Comments on the Global Structure of Instructional Materials
6. Lines of Development for the Major New Industry of Multimedia Electronic Materials Production

1. Introduction

The continuing increase of computer speeds and the falling costs of computer processing and data storage are creating remarkable new possibilities for enriching the informational environment of daily life, through a fusion of microcomputers and electronic home entertainment devices. To envision what is becoming possible, imagine a computerized environment with a large wrap-around screen on which one can display many high-resolution images, both computer-generated and retrieved from videodisk, simultaneously or separately, in various sizes and tempos. Imagine that it is possible to access all of the world's books (scientific treatises, journals, novels, works of criticism, archives) within this environment, in a manner which makes elaborate indexes, cross-references, and illustrations instantly available. Suppose that images of all the world's great works of art are also available, along with images of all notable buildings, landscapes, and other scenes, the whole providing a 'universal museum' of art, architecture, culture, and landscape. Imagine that, within the virtual concert and performance halls of this universal museum, all the world's concerts, operas, and dramas are continually in performance. Imagine that the museum visitor, seated at a console which gives unlimited access to this universal library and museum, can soar freely through any part of it simply by moving a joystick; can summon up any work of literature just by touching an on-screen reference to it; can summon up any existing audio or video dramatization of a literary work; can view any photograph or film relevant to a scientific or historical text, simply by touching the screen to indicate what is desired. Imagine that moderately personalized cost-free tutoring is available to any student wishing to master some part of this immense edifice of knowledge and information, which is, in fact, a matrix into which much of human culture can be drawn. Imagine finally that artists and scientists are continually making additions to the basic content of this universal library and museum, and that scholars, critics, and educators add continually to the fabric of crossreference, comparison, and explanation which ties it all together.

2. Technologies

The present and prospective technologies basic to these dazzling possibilities are as follows. Computation rates of 10 MIPS (millions of instructions per second) will soon be available on a single chip. Current laser videodisks store 54,000 frames of video information, each frame potentially amounting to (1/2) * 512 * 512 * 2 bytes (i.e. 256 Kbytes) of information, for a potential of roughly $13 * 10^9$ bytes of information per disk. Considered as stored text at 4,000 characters per page, such a disk could hold approximately $3 * 10^6$ pages, or 6,000 books of 500 pages each. This information could also be played back at the approximate rate of 50Kbyte/sec characteristic for digital audio, rather than the 8 Mbyte/sec required for video. The resulting expansion factor of roughly 160 would make it possible for 80 hours of recorded sound to be held on a single disk in the same space occupied by 30 minutes of video.

These already large storage capacities can be enlarged further by use of statistical encoding techniques. Huffman-coding of text affords roughly a 4:1 advantage, potentially allowing a single disk to hold 24,000 volumes of printed information. Recorded audio can be represented by Huffman-coded deviations from a linear (or more sophisticated) prediction of wave-intensity values, gaining a 4:1 storage density advantage. This might allow approximately 640 hours of audio to be stored on one disk. Two-dimensional analogs of this prediction scheme can be used for still-frame video. For example, a video frame can be separated into its three constituent colors, and for each color the average of the intensity value taken from the pixel U immediately above a pixel P and the intensity value taken from the pixel L immediately to the left of P can be used as the predicted intensity of P. This simple scheme should attain 4:1 compression, and might make it possible to store 220,000 still images on a disk. For moving video, P's value can be predicted as the average of the intensities of U, L, and of the intensity of P in the immediately preceding frame. Assuming that this technique attains an 8:1 compression ratio, it might become possible to store 4 hours rather than 30 minutes of moving video on a disk.

These capacities can be enlarged still further, and by very large factors, by using 'animations' rather than simple 'excerpts', i.e. by generating the audio and video information dynamically from a higher level symbolic representation, rather than simply playing back from a recording. Music, for example, can be represented by coded scores for a very sophisticated synthesizer. In such scores, each 'note' would indicate frequency, intensity, and duration, as well as 'attack' and 'decay' and the overtone structure required to produce any desired 'instrument sound'. At 10 bytes/note, 5 notes/second, 50 instruments per ensemble, this requires 2.5 KB/second, potentially allowing one disk to represent 3,200 hours of very high quality synthesized music.

Corresponding possibilities exist for video. Suppose, for example, that we record very high resolution images of paintings, e.g. images formed at 8000 * 8000 resolution. Assuming that a 2-D predictive encoding can gain an 8:1 storage density advantage for images of this sort, roughly 50 video frames would be required to store a single ultra-high quality image, and thus 1,000 such images could be held on a disk. Suitable sampling and reconstruction hardware, operating in realtime, could then reconstruct views of the image from a dynamically varying viewpoint.

A more advanced technology of this kind might allow arbitrary views of 3-D objects such as sculpture to be reconstructed in realtime from prestored stereo images. Such a technology might allow a single disk to hold all the information needed to reconstruct views of up to 500 solid objects, urban scenes, architectural interiors, etc.

Many of the devices necessary to realize such multimedia electronic environments already exist, though not always in the ideal forms which they are apt to evolve. Aside from basic computational and storage devices, the list of desirable devices includes:

(i) A 'surround screen' or multiscreen, on which multiple dynamic, high resolution images can be presented at all points of a viewer's visual field;

(ii) High-fidelity digital audio output, with variable apparent source;

(iii) High-quality computer-generated voice;

(iv) Voice recognition, at least for input of simple commands;

(v) Touchscreen;

(vi) Electronic audio output devices, e.g. high quality digitally-sampled audio synthesizers with multi-instrument/voice capabilities;

(vii) Powerful realtime computer graphic capabilities, allowing decoded video images to be fused with graphic images, text overlays, and moving images produced by animation of high-resolution natural-object still images. These fused images should be displayable in arbitrary sizes and combinations.

The image display capabilities of present videodisk-based systems can easily be improved by furnishing the disk-player with multiple read-heads (e.g. 8) so that stills, moving images, audio and text information can all be read simultaneously from a single disk, thus preventing unwanted pauses during searches. Writable magnetic or optical storage systems of a few gigabytes capacity are also desirable for use as temporary storage to accelerate access to the images, sound, or text required by any particular educational or artistic end product. A somewhat more advanced storage technology might be able to provide removable 'disk-sets', e.g. of 25 optical disks, easily attachable to and detachable from a common spindle. This could enlarge the amount of material accessible at any one time to 600,000 volumes of text, or 80,000 hours of synthesized music, 8,000 hours of recorded sound, 5 million still frames, or 100 hours of video.

3. The Computerized Electronic Environment as a Medium

The basic technology reviewed in the preceding section provides the underpinnings for an astonishing new artistic and educational medium. The inspired and sustained efforts of gifted artists, educators, and scholars will be needed to explore the potential of this new medium adequately. Nevertheless, it may be helpful to attempt an outline of the basic elements which such a technology can put at the disposal of a creative artist/educator, for integration into higher-level educational structures such as those described in the next section. The following headings identify some of these elements:

(i) An *excerpt* is a particular segment of text, audio, or video, displayed with or without motion, for a fixed or variable interval of time. Examples include a screenful of information presented as text, a sequence of phrases produced by a computer-generated voice or spoken by a recorded voice; a musical passage; a still-frame displayed as an illustration; a video passage played back from videodisk as a moving illustration.

(ii) A *parametrized excerpt* is an excerpt dependent on parameters that can be varied to achieve educational or aesthetic effects. Examples are: text which can be visually highlighted in a variety of ways; computer generated voice with variable stress, rate, intonation, and timbre; images which can be viewed in a variety of lightings or from a controllable vantage point.

(iii) An *animation* is a dynamic (and thus inherently modifiable) excerpt generated by continuous or discrete variation of the parameters of a parametrizable excerpt. Examples are: moving images of works of art generated from high-resolution still images by varying the point from which they are nominally viewed; 'sound-and-light' effects generated by presenting still images under varied illuminations; and perhaps music itself, if music is viewed as the result of 'animating' a static score.

(iv) A *mix* or *collage* is a collection of excerpts, displayed together for combined effect. Examples are: graphic overlays on video images; video combined with music; text combined with still-frame or moving illustrations; juxtaposition of several graphic and video images having deliberately varied sizes; crossfaded combination of multiple voices, possibly with music, to achieve dramatic or choral effects; combination of text with voice or music for emphasis.

(v) A *parametrized mix* is a mix whose constituent elements and element interrelationships can be varied by changing one or more parameters. Examples are: simultaneous video presentations of multiple images whose relative sizes can be varied to create animation-like effects; audio mixes in which the relative intensity of various sound elements are varied to train a listener to perceive sounds that might not ordinarily be noticed (e.g., training novice physicians to detect various types of heart murmurs.)

Mixes and parametrized mixes allow indefinitely many excerpts and animations to be produced from an initial stock of elements, e.g. allow 'recycling' of art fragments for a great variety of purposes and effects. Examples in the humanities might include the use of images of historic locales for sound-and-light recreation of historic events; or combined use of multiple images and of text for comparison of artistic styles across periods or national or regional schools. Examples in science are the use of a natural scene as the starting point for an extensive lesson in biology and ecology; use of multiple tissue images, acquired at a variety of macroscopic and microscopic scales and combined with computer-generated graphics to illustrate the external appearance, histology, and physiology of a disease; overlay of explanatory graphics on the filmed version of a physical experiment which relate observed phenomena to underlying principle.

4. Educational Use of the Multimedia Environment

To convey some idea of the way in which the techniques described in the previous section can be used, we begin by presenting several hypothetical lesson scenarios. Subsequently we reflect in more general terms on the basic content of the educational process and the manner in which multimedia computerized systems can accelerate learning.

Scenario One: *A Lesson in Ancient History.* The student, in a college-level world history course, is reviewing the sequence of political events surrounding the collapse of the Roman Republic and the rise of the Principate of Augustus. Some of the material is presented as straight text, which appears on the screen in an attractive book-like format, with still illustrations occupying roughly a quarter of the screen. This text material can be enlivened by occasional (optional) audio quotations from ancient sources, such as Plutarch's *Lives of the Noble Romans.* Quotations from each such source can be read in the voice of a different actor who represents the ancient author. This material is announced on the screen by a line reading something like: "Concerning Crassus, Plutarch said the following...". This line appears in a color which the student has learned to recognize as meaning that an optional audio insert is available. If he touches this line on the screen, the video illustration on the page switches to show a bust of Crassus, while Plutarch's remark on Crassus is read in a voice emphasizing its drama and irony. At a later point in the lesson, information concerning the Battle of Philippi is presented, first by text explaining its date and significance, etc., but then as 5-10 minutes of filmed material taken from videodisk and introduced by a compelling musical fanfare. A map of Italy and Greece, showing the lines of march of the Republican and Antonian armies to the battle, is displayed. This is followed by a more detailed map of the area of northern Greece in which the battle was fought; then by video footage, taken from a helicopter, of the terrain of Philippi. After this, a brief enactment of the battle itself is shown, with views of its commanders in action, voice-over commentary on its main incidents, and dramatic scenes of the suicides of Cassius and Brutus. These video-clips then yield to text reviewing the significance of the battle and the events which followed it, with audio statements from ancient authorities reflecting on its meaning for the Roman world.

This presentation is followed by an interactive multiple-choice quiz on the material of the lesson. Questions answered incorrectly trigger re-runs of relevant portions of text, audio, or captioned stills. The student can also make abbreviated extracts of the text, audio, and video material presented, in order to build up a personalized course summary for end-of-term review.

Scenario Two: *An Arithmetic Lesson.* Here, we envisage a grade-school student in the process of mastering "long" addition (addition of multi-digit numbers). The student is assumed to have memorized the single-digit addition table, but the lesson software remains alert for errors in single-digit arithmetic, and will review items of this material, using appropriate displays and audio statements, when appropriate. It is also assumed that, in a prior lesson, the student has practiced distinguishing sums like 6 + 5, which yield a result of 10 or more, from sums, like 4 + 3, which add to less than 10. Finally, it is assumed that the idea of 'carrying' basic to multi-digit addition has just been presented in a classroom lesson.

The computerized lesson begins with a brief video clip of a teacher reviewing this basic idea at a blackboard; the teacher works several sums, points to relevant features of the numbers being added, writes ‘carries’ on the board, emphasizes the right-to-left nature of the addition process, etc. Then the student enters an interactive review and practice sequence covering the material just presented. A practice sum, for example,

```
    375
+   419
---------
```

appears on the screen. A voice prompt is given, telling the student to participate in adding these numbers. However, the student is only required to handle the carries; the computer will handle the single digit additions automatically. As the turn of each column to be added arrives, the column is highlighted by color, blinking, and a graphic arrow. A two-item selection menu, offering a choice of "carry" and "don't carry" appears on the screen. The student is expected to touch the appropriate one of these, causing the computer to write the associated sum-digit in the appropriate column and advance to the next column. Any error will trigger appropriate video and audio diagnostics, plus a short review of earlier material.

A first illustrative sum is run through in completely automatic fashion, with a graphic hand representing the student’s finger, and an audio commentary explaining what is expected at each point. Then several interactive examples requiring student response are presented. If these are handled successfully, the student can be offered an opportunity to play a video game which tests the speed with which she is able to handle such problems. In this game, moving spaceships bearing long sums appear on the screen, under pursuit. The student continues to interact by touching the carry/don’t-carry menu. Incorrect responses cause the spaceships to explode, losing a point; correct responses allow them to escape. Delayed responses cause the spaceships to slow down, making them vulnerable to demolition by their pursuers.

After this lesson subsegment, drill resumes. Now the computer handles the carries, and the student is required to supply the associated sum-digit. This is done by touching a 10-item menu of digits on the screen. Next, the student is asked to handle both sums and carries together, and then finally to handle digit-by-digit long addition without any computer assistance. Errors indicating problems in handling either carries or single-digit addition trigger reviews of prior material. Review games at all skill levels are available.

The arithmetic-lesson scenario just presented illustrates several principles important to courseware design.

(a) Ways must be found to focus students’ attention on the skills being taught, not on ancillary details of the computer environment. For example, use of a typewriter keyboard distracts attention from material on the screen, and should be avoided whenever possible, in favor of touch-sensitive menus displayed directly on the screen. Similarly, having to read written instructions distracts from other screen items, and should be avoided whenever spoken instructions or graphic emphases can be used.

(b) Many conventional academic skills simply amount to the ability to select and apply algorithmic or near-algorithmic procedures rapidly and correctly. These skills are built up from subskills, e.g. to add long digits one must know that the process moves from right to left, must handle single digit sums correctly, and must know how to manipulate carries. Computer-controlled interactive courseware can concentrate on one subskill at a time, in a manner impossible for a textbook and hardly available to the classroom teacher, namely by asking the student to handle only that part of a procedure on which pedagogical stress is to be laid, while other aspects of the same procedure are handled automatically by the computer.

Scenario Three: *A Geometry Lesson*. The geometry-lesson scenario which we now present emphasizes this last point, but in a more sophisticated way than the simpler arithmetic lesson. It also shows the extent to which CAI lessons will have to be backed up by complex symbol-manipulating software. We envision a high-school level student working through a conventional plane geometry course which emphasizes proofs and the solution of "original" problems. The basis of many elementary proofs of this kind is the systematic use of triangle congruence to prove equality of various angles and distances in a given figure. A computer can easily apply these rules systematically and rapidly: a proof can therefore be considered entirely elementary if it can be achieved using the triangle congruence method repeatedly, and if no other idea is required. However, if this one tool is insufficient, a less elementary proof must be given. Often the key to such a proof is to add a few lines or curves (e.g. circles) to the original figure: the proof is then elementary in the expanded figure.

An interactive lesson can exploit this fact, which is crucial to conceptual mastery of high-school geometry. For this, a succession of problems requiring proof, each accompanied by a figure presented on the computer's display screen, can be presented. The computer can then reason as far as possible using congruent triangles only. It can list its conclusions, and, if the student so requests, can review the steps of inference which led to each conclusion. (Such review will emphasize the figure elements involved in each step graphically, and will use audio commentary. If requested by the student, statements of any prior theorems used in the computer's inferences can be retrieved and displayed or spoken.) However, for problems too difficult to yield to such elementary reasoning, the student will have to discover the line (or lines or curves) which will make elementary proof possible when added to the figure. As a hint, the computer offers a multi-item touch-sensitive menu: "Shall I add a line?", "Shall I add a circle?", etc. Suppose, for example, that the student choses to add a line. Then a menu of possible lines can be offered: "line between two points", "angle bisector", "line perpendicular to specified line through specified point", etc. When the student choses one of these by touching it, labels can appear in the figure being analyzed, and then by touching these labels the student can signify exactly what new item should be added to the figure; these elements will appear immediately in the figure. (They can be erased just as easily if the student decides that they are not wanted.) Once these additions are complete, the computer will apply elementary reasoning-by-congruence to the expanded figure. If the proper lines or curves have been added, the proof will now succeed; if not, the student must continue wrestling with the problem.

This scheme, which combines student interaction with sophisticated computer assistance, has the merit of focusing attention on the key strategic and conceptual decisions needed to handle a problem, while lower-level, assumed skills are automatically supplied. This recipe is applicable to many other situations in mathematics and science. For example, a student learning to deal with word problems in algebra may only be required to read English-language problem statements and convert them to algebraic equations; the computer can then solve these algebraic equations automatically. Similarly, a calculus student learning how to evaluate indefinite integrals can be offered a menu of standard manipulations: "integrate by parts", "substitute for a variable", "decompose the integral as a sum", "add and subtract a term to the integrand", "factor", etc. Once an alternative has been chosen and its parameters supplied, the necessary manipulations can be performed automatically. As already said, by easing the learner's burden of low-level arithmetic and algebraic details, we can focus attention on the conceptual content of the material on hand, thus permitting deeper comprehension. This should help both the strong student, whose progress the computer can accelerate, and the weaker student, for whose inaccuracy in applying the auxiliary manipulations which support higher level skills the computer can compensate. Hopefully, as weaker students come to understand how their skill deficiencies impede their ability to handle interesting problems, and as interactive experience pinpoints these deficiencies, they will be motivated to overcome their own shortcomings by systematic review and practice.

Many other equally compelling scenarios illustrating the combined use of video, sound, and computer graphics could be presented. In courses such as physics or chemistry which traditionally involve laboratory demonstrations, video clips of demonstration experiments, highlighted appropriately by dynamically evolving graphs and charts, can be presented. Significant aspects of a demonstration can be emphasized using audio sound effects, e.g. the notion of acceleration in elementary physics can be emphasized by emitting clicks at a rate proportional to the velocity of an accelerating body. When quizzes on the significance of an experiment answered incorrectly, specially highlighted video reruns of key moments in the experiment can be called up. Experiments in genetics or biology can be presented in time-lapse video representations to emphasize principles of growth. Video clips can display a much larger variety of experimental equipment than could conveniently be shown in a university laboratory.

Simulations can also be used very effectively. For example, a student in an organic chemistry course could be asked to design molecules, which he specifies by connecting carbon, hydrogen, nitrogen, and other atoms in various patterns using a graphic editor. The actual properties of the molecules suggested can then be retrieved from an on-disk chemical encyclopedia and compared to the desired properties. Various 3-D representations of these molecules could also be displayed, allowing the student to develop further intuitive understanding of their chemistry.

Similar uses of videodisk for social science training are possible. For example, extended psychotherapeutic interviews with a series of cooperating patients can be segmented into question-response pairs and stored on videodisk. Student therapists can then be presented with a menu of possible interview questions and asked to decide what sort of response they would expect such a question to elicit, following which the patient's actual response can be retrieved and played back.

These scenarios exemplify our expectations concerning educational use of an electronic environment. Having presented them, we now attempt to reflect in more abstract terms upon the basic content of the educational process and the manner in which multimedia computerized systems can support and enhance this process. What is it that one learns when one does learn? Though no finite statement can capture all aspects of this process, which is as elusive as human psychology itself, perhaps a few abstract phrases can establish a useful conceptual framework. Learning can be regarded as the acquisition of *facts* (and their interrelationships); *procedures* (for reaching significant goals); *methods* out of which procedures are composed; and *criteria* for the applicability of particular methods. To reach higher levels of creativity, one must become aware of many *examples* illustrating the use of specific methods and procedures, and come to perceive *analogies* between the situations in which particular methods can be applied, and between methods themselves. The multimedia electronic environment which we have portrayed can be a powerful aid for conveying all these basic elements of education.

In regard to facts, its capabilities are superb. As already indicated, an electronic environment can make an exhaustive mass of facts (eventually, all the world's books) available, and can use extensive audio, graphic, and visual materials to assist comprehension of these facts. The universal library and museum provided by an electronic environment can be available in fully cross-referenced, readily 'browsable' form. Emphasis and drama can be added in a great variety of styles to highlight significance. Sophisticated summaries, cross-indices, and semantic links, not only connecting text to text, but also linking text to video and audio, video to text, text to simulations, can all be constantly available. At any point in any of the textbooks, journals, novels, or treatises available within the universal electronic library, screen touch can be used to access extensive indices listing every significant term that appears in the text, as well as indefinitely many supplementary volumes of more detailed information accessible through the same index. As an example of what this means, consider again the student of history (or casual reader of a historical novel) whose exploration starts with a reference to

Louis XII..

but who then at successive moments and touches can summon up

..biography, Encyclopedia Brittanica..
..artworks of period of Louis XII on view in Museum..
..major composers of periods, trigger performances if desired..
..wars of..
..role of Louis XII in development of the French Monarchy..

For another example, imagine a computerized 'nature walk' which begins with a moving image of a hillside path along which various trees and shrubs are seen, and along which the presence of various animals is evident through sounds (such as birdcalls), fleeting glimpses, tracks, and openings of burrows. By using the touch screen to indicate interest, the browser could then access detailed closeups of fauna, on the ground or in their subterranean habitats, encyclopedic information concerning plants, including time-lapse images of their diurnal and seasonal cycles.

It is also worth noting that the system user's browsing activities can be supervised by a programmed *exploration supervisor*, which can maintain a history of all informational nodes traversed by a user, thereby facilitating backtracking and exploration of collateral branches. This same *exploration supervisor* can collect information on a user's interests, and can employ this information to optimize the order and selection of items presented subsequently. It is also easy to allow definition of personalized information subcollections ('favorite' art items, or items having special significance to a student preparing for an examination, or to a professional who needs a 'desktop summary' of some more extensive mass of information).

Exploration and concentrated study can both be assisted by providing various standard semantic (as distinct from 'index' or 'keyword') links in the universe of facts available within an electronic environment. Among such links, which one expects to be able to traverse at a touch, one would expect to find the following:

WHY? Justify an assertion made in a text.
WHAT? Explain a term or concept occurring in a text, or an element appearing in a graphic display or video image.
SO WHAT? Explain the importance of a fact, concept, or method.
HOW? Explain or illustrate a method or activity mentioned in a text.
CRITICIZE. Provide material refuting an assertion or denying its importance.
SHOW. Display pictures or audio representations of items mentioned.

As an example of this last form of semantic link, imagine a student leafing through a history of U.S. Civil war who encounters the remark that "Shortly after this crucial battle, Lincoln visited its site and delivered a brief memorial address which has become one of the most famous in American oratory..." Touch-invocation of a semantic SHOW link to this sentence could then produce a spoken dramatic recreation of Lincoln's Gettysburg address.

DESCRIBE OTHER SIMILAR THINGS. This semantic link should summon up facts and concepts generalizing, specializing, or standing in significant analogy to an item mentioned in text or appearing in the universal museum or one of its attached concert halls.

So much for the treatment of *facts* within the electronic environment. In regard to the teaching of *procedures*, we note that procedures are normally learned together with batteries of more general *methods*, using which specific, particularly significant, procedures are composed. For example, the student of cooking learns both complete procedures for producing specific dishes (e.g. 'Baklava' or 'Szechuan-style dry-fried pork') and underlying methods such as mincing, fine slicing, dicing, stir-frying, hot-oil crisping, characteristics of generally useful ingredients such as philo, broth, ginger, hot peppers, cornstarch, etc. The student of algebra learns not only specific formulae but also general techniques such as substitution, transposition, graphic estimation of root position, etc. Since procedures generally exist within a larger environment of applicable methods (and since they must often be applied to partly unpredictable and refractory materials), an essential part of a craftsmanlike knowledge of procedures is the ability to recognize situations in which particular methods should be used as elements of a composite procedure; also the ability to recognize positive signs that a method is having a desired effect or danger signals which indicate that it is starting to fail; and finally knowledge of methods for recovering from situations in which some step of a procedure (perhaps a step applied erroneously) has produced some unwanted effect. Moreover, if complex, initially invisible internal states control the behavior of an object being manipulated, one must learn to think in terms of those states, and must acquire diagnostic procedures for determining what the state of a system is whenever this is unknown and important. In technical environments in which procedures are applied in partly automated fashion, mastery of subprocedures for calculating the setup parameters which control automated procedure runs (e.g. baking times and temperatures) is an important aspect of procedural learning. Computerized systems can aid in this activity by providing computational help in determining required parameter values, by displaying the consequences of suboptimal parameter setups, etc.

These general remarks apply with equal force to electronic equipment repair, to the baking of breads and cakes, and to the acquisition of expert knowledge in physics, mathematics, and chemistry.

Used artfully, the computerized multimedia environment can accelerate the acquisition of all these aspects of procedural knowledge. First of all, it allows procedures to be studied in a simulated environment over which the learner can be given relatively complete control. Systematic comprehension of the battery of applicable methods can be aided by displaying menus of methods; perhaps just those methods actually applying to the particular stage which a procedural simulation has just reached, perhaps a longer list also including other methods which the student might (wrongly) imagine to be applicable. Images of the state reached by a system being manipulated procedurally can be displayed as the procedure progresses. The display of information regarding the state can include both ordinary macroscopic information and normally invisible microscopic information where such information is significant or controlling.

Timescales can be accelerated to allow rapid exploration of phenomena which normally involve delay due to some requirement for protracted 'ripening', or can be slowed down to allow observation of phenomena which normally proceed too rapidly for comprehension.

As an example of this, consider the way in which a computerized multimedia environment could be used to teach the art of baking. Here one proceeds from initial ingredients which may have to be carefully chosen (high-gluten flour behaves differently from very starchy flour) or specially prepared (sifted, perhaps), through a series of intermediate stages (mixing, stirring, folding, rising...) which have interesting and crucial rheological effects, to final stages of baking and decoration. This is an art requiring time to master, particularly because the number of experiments that one can perform in a single day is limited. The electronic environment can ease and accelerate this process in many ways. The characteristic external appearance of flours and doughs of various grades and qualities can be displayed, together with their internal rheology, using a combination of macroscopic images, microscopic views, and schematized graphics, for example the manner in which dough reacts to pounding or to pulling can be correlated with the average length and orientation of the gluten strands within the dough, proportion of oil or saturated fat, etc. All this can in turn be related to flow quality, mixing proportions, and amount of hand or food-processor kneading. Techniques can be explained and illustrated for determining when dough to be used for one purpose or another has been kneaded sufficiently or excessively, contains too much or not enough oil or water, etc. Real or graphic cross-sectional images can be used to explain the process of rising, as driven by yeast or baking powder, and to clarify the manner in which the average 'gassiness' of a lump of dough can be deduced from its external appearance and reaction to various standard baker's tests. The reaction of doughs of various consistencies and thicknesses to baking at various temperatures and for various periods can be shown by time-lapse photography both of loaf externals and cross-sections, and by graphics describing salient aspects of internal molecular state. Techniques for creating special qualities (crustiness, chewiness, or crispness) or special geometries (pretzels or braided loaves) can be shown. All this information can be presented in an environment which allows the student to make mistakes and observe their consequences, backtrack to the point of an error and try something else, vary proportions or processing times to achieve special results. Recipes and comments by master bakers can be available for consultation at any point. If the procedure-exploration environment that we have been describing is well-constructed, it should enable its user to gain understanding of many of the aspects of baking much more rapidly than would be possible in an actual bakery, in part by greatly compressing the time needed to apply various essential processes.

It should be clear to the reader that this example can easily be modified to apply to baking-like processes in many areas of science and engineering, such as chemical and medical laboratory procedures, processing of silicon wafers, and metallurgical processes. The basic idea of using visual/graphic/textual exploration and observation to find a specific procedural path within a tree-structured virtual environment of methods and intermediate states is applicable to a very wide range of practical and intellectual activities.

As noted in our earlier geometry lesson scenario, much of the difficulty of learning procedures which are more intellectually demanding, (e.g. symbolic integration in calculus, or solution of small systems of algebraic equations by formation of resultants) comes from the fact that the procedure to be learned requires use of various subprocedures (e.g. differentiation, trigonometric substitution, algebraic simplification, and determinant evaluation) which the student attempting to learn a new skill may not have mastered fully, so that struggles with these subskills may distract and erode confidence. A computerized system can deal very effectively with this common pedagogical difficulty, focusing the student's cognitive attention by performing subprocedures automatically. One can focus attention on the battery of new methods to be acquired by displaying menus of methods and by allowing the student to compose programs which apply those methods automatically, in sensible sequence. The arithmetic and geometry lesson scenarios presented earlier illustrate these ideas.

Criteria for determining when particular methods should or should not be employed can be taught by showing the typical physical appearance or abstract structure of materials or situations which call for application of these methods. *Examples* can be conveyed by displaying many of them and by visual or graphically highlighted comparisons, and *analogies* can be conveyed in the same way or by using coupled animations which emphasize abstract similarities between the workings of related methods.

5. Some Additional Comments on the Global Structure of Instructional Materials

The overall structure natural to a body of instructional materials will of course depend on the conceptual substance which it aims to present. We have already described the kind of 'exploration' environment likely to typify presentation of extensive bodies of fact. As already noted, this can be expected to be highly interactive and densely interconnected via semantic links and links carried by indexed keywords. Lessons which convey procedural knowledge can be expected to reflect the graph-like abstract structures of the procedures that they teach. Aside from these two important special cases, instructional material is apt to have a looser, network-like, overall structure, composed of individual information-presenting nodes connected by a mesh of menu-like choices, either explicit (as in actual menus) or implicit (as when lesson software acts to classify a system user's answer to a question into one of a finite collection of possible bins, each of which advances him from the environment in which the problem is presented to one of several possible successor environments.) >From the global point of view, the question/response paradigm acts just like a menu-choice operation, the difference being only that the 'menu' involved is concealed from the system user.

As its user advances between nodes of an instructional network, the software system with which he is interacting can collect information concerning his behavior, preferences, and progress. To the extent that it proves possible to analyze this (often enigmatic) information, the system can build up a model of the student which it can use to modify the excerpts, animations, and simulations presented at each informational node subsequently traversed. If available, such a model can also be used to control the way in which he is routed through the network, e.g. signs of difficulty can trigger presentation of auxiliary material which is bypassed by students showing signs of more rapid mastery.

6. Lines of Development for the Major New Industry of Multimedia Electronic Materials Production

Developers of truly effective computerized commercial and educational systems will draw upon many sources of ideas, ranging from recent developments in cognitive science to the informal but essential wisdom of the skilled teacher. However, the core of the present opportunity is technological and developmental: to build a hardware/software environment within which attractive, exciting, carefully directed, smoothly flowing, and richly informative material can be brought to the individual viewer/auditor; to create those initial demonstrations from which a broader commercial activity can grow; and to assess/refine this experimental material by observing the way in which typical users interact with it. A necessary initial step is organization of several production facilities within which multimedia electronic workstations, supported by software systems that facilitate their use, are available. A core staff will be needed to provide first-class hardware, software, graphics, and image-processing support to teams of subject-matter specialists, experienced teachers, and visual artists with creative ideas for composing individual presentations. The key persons around whom the activities of such 'studios' are likely to revolve will be 'director-like' figures combining integrative and organizational skills with scholarly and artistic interests and abilities.

The rate at which the envisioned electronic multimedia environments, and attractive courseware and entertainment materials for them, can move from small-scale laboratory use to general public acceptance and broad involvement by scholars, educators, and artists will depend on the sequence in which profitable business opportunities surrounding this development materialize and are taken up. Although the necessary technological developments are already proceeding very rapidly, one crucial link, wide availability of standardized systems on which expensive multimedia materials can be delivered to audiences of millions, remains to be filled in before the production of such materials can attain film-industry proportions. This will not be an easy gap to fill, since it poses a chicken-and-egg problem: without standardized delivery systems, no broad audience, hence few high-quality materials; without attractive materials, little incentive to purchase, and hence less incentive to manufacture, the delivery systems that are wanted. This knot can most readily be cut by a corporation powerful enough to drive the market, or by government involvement, without which these developments may proceed only haltingly.

Nevertheless there exist smaller-scale opportunities which should serve to push things along. CD/ROM technology is rapidly being developed for home use and should shortly become a standard delivery vehicle for numerous encyclopedias and other text databases for which first-class indexing is strongly advantageous, e.g. for major collections of scientific abstracts, back-series of archival journals, product and manufacturer catalogs such as the well known Thomas Register, etc. Immediately profitable opportunities also exist in publication of training materials for 'salaried students', e.g. corporation employees (such as airline pilots or bank tellers) who must receive training at employers' expense. Similar opportunities exist in production of military training materials.

A second immediately profitable possibility lies in the development of a line of videodisk-based *Commercial Kiosks*. These can be somewhat reduced versions of the deluxe multimedia environments that we have described. Such Kiosks can be used very effectively to deliver information to customers in large department stores, discount houses, and other commercial enterprises with goods and services to advertise. Information presented in video, audio, and text formats can include location of departments and services, images of items on sale (with ads), daily or weekly specials, catalogs of additional items or services available by mail order, etc. These same Commercial Kiosks could also support direct credit-card authorized interactive purchasing. One can expect much of the cost of video materials production to be borne by the manufacturers of featured products, who will presumably be as eager to compete for on-disk display space as they presently are to compete for shelf and rack space. Another interesting possibility is that of using interactive Kiosks aggressively by installing sidewalk Kiosks that would alert shoppers about to enter a competing shop about bargains or superior product lines available at one's own shop.

Once Commercial Kiosks of the kind we project came into widespread use, one could expect the videodisks used on them to become available for home viewing; this would create an electronic version of the familiar 'Sears & Roebuck' sort of catalog merchandising.

The travel industry offers additional opportunities for early introduction of delivery systems similar to the proposed Commercial Kiosks, and of home equivalents of such systems. Once very high-resolution displays and images of the sort previously described become available, travel agents should find it advantageous to let prospective clients examine scenes which they were thinking of visiting, view hotel and casino interiors, inspect museum and palace settings, beaches, etc. Persons planning tours might find semi-instructional disks of the 'Visit Paris and Learn French' sort very appealing, especially if these included extensive guidebook information concerning inexpensive hotels, train schedules, driving distances, etc. Tourism promotion agencies, major airlines, etc., should all be willing to subsidize production of such materials.

Commercial and training activities of the sort just described will require, and hence should spur development of, a variety of technical instruments, generally not going very far beyond what is already available, but somewhat more highly functional, better integrated, and substantially easier to apply in commercial environments. These items fall into two principal classes: relatively inexpensive equipment intended for use in commercial, school, and home information delivery stations, and substantially more elaborate and expensive items of studio equipment, to be used for production of commercial-grade audio and video materials.

The equipment components wanted for the delivery systems themselves have already been discussed sufficiently. Aside from basic computing engines and storage devices, the list of studio-related items likely to be in commercial demand includes:

(i) Audio and video compressors and encoders;

(ii) Video mixers and special effects hardware; digital animation stands;

(iii) High speed hardware for composition of computer graphics and for manipulation of natural images, especially '3-D images';

(iv) Hardware for acquisition and encoding of natural images;

(v) Very high capacity storage systems for easy access to major image archives;

(vi) Videodisk writing hardware.

On the software side, much improved authoring languages will certainly be required. These will have to support composition and maintenance of very large educational and other information-delivery systems. Besides such basic capabilities as first-class response analysis primitives, these languages must provide powerful facilities for creating and reviewing large lesson structures, and facilities for defining dynamic multimedia effects such as animations. Encyclopedic browsing capabilities will require various forms of specialized software, including indexing and database software able to add masses of information to still larger pre-existing systems. ===========

T

his book may be kept

FOURTEEN DAYS

A

fine will be charged for each day the book is kept overtime.

V. 17 1988

MAY 10 1991

JUN 07 1991

FEB

YLORD 142

PRINTED IN U S A

GA

NYU COMPSCI TR-211
Schwartz, Jacob T
'Dreamworld' c.2

NYU COMPSCI TR-211
Schwartz, Jacob T
'Dreamworld' c.2

DA

MAY 10 1991 (SAKOLITZ?

~~JUN 07 1991 RENEWED~~

[illegible] Z. Liu

LIBRARY
N.Y.U. Courant Institute of
Mathematical Sciences
251 Mercer St.
New York, N. Y. 10012